

Modeling Spacecraft Safe Mode Events

Travis Imken
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109
818-354-5608
Travis.Imken@jpl.nasa.gov

Michael DiNicola
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109
818-354-2558
Michael.Dinicola@jpl.nasa.gov

Thomas Randolph
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109
818-354-4871
Thomas.M.Randolph@jpl.nasa.gov

Austin Nicholas
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109
818-354-5608
Austin.K.Nicholas@jpl.nasa.gov

Abstract — Spacecraft enter a ‘safe mode’ to protect the vehicle when a potentially harmful anomaly occurs. This minimally-functioning state isolates faults, establishes contact with Earth, and orients the vehicle into a power positive attitude until operators intervene. Though ‘safings’ are inherently unpredictable, mission teams build in time margin during operations to determine root causes and restore functionality. Planning and managing this margin is both critical and enabling on mission architectures dependent on near-continuous operability – such as a low-thrust electric propulsion mission.

To better quantify the occurrences and severity of safe mode anomalies, the Jet Propulsion Laboratory (JPL) has assembled a database of safings from past and active missions. Currently nearly 240 records are captured from 21 beyond-Earth missions, stemming from a collaboration between teams at JPL, Ames Research Center, Goddard Space Flight Center, and the Johns Hopkins University Applied Physics Laboratory. This paper discusses the event database, explores a statistical approach in modeling the occurrences and severity of safing events, presents a simulation technique, and details recommendations and future work to benefit future concepts.

TABLE OF CONTENTS

1. INTRODUCTION AND MOTIVATION	1
2. SAFE MODE EVENT DATABASE	2
3. MODELING OCCURRENCES AND RECOVERIES ...	3
4. APPLICABILITY TO FUTURE MISSIONS	7
5. RECOMMENDATIONS AND FORWARD WORK	9
6. CONCLUSIONS	10
ACKNOWLEDGEMENTS	11
REFERENCES	11
BIOGRAPHY	11

1. INTRODUCTION AND MOTIVATION

Since the development of the Galileo mission¹, spacecraft designers have implemented an onboard, self-detecting

failsafe mode. These ‘safe modes’ are architected to guarantee a vehicle can wait in a safe operating state until ground operators can recover and restore nominal operations. The characteristics of this safe state are common across missions: isolate faults to prevent further propagation, establish and maintain communications with Earth, orient the vehicle in a power-positive and thermally stable configuration, and be able to maintain this state indefinitely. Though some in-flight safe mode entries are planned, such as a reboot to initialize a flight software update, the majority of ‘safings’ are unpredicted events or conditions that trigger a system-level fault protection response.

This research considers the subset of anomalies that result in safe mode entries. Spacecraft fault management is designed to handle a spectrum of faults and anomalies seen across software, hardware, and payloads. When an anomaly is detected, fault protection will first work to isolate the issue at the component level. If the fault persists, subsystems or instruments can be isolated and marked as ‘sick’. At this tier, the issue will be noted in the next communications pass and the vehicle will continue with nominal operations. However, further fault containment failures will trigger a system-level response, often leading to a safe mode entry. Safing events stand apart from other flight anomalies because of the technical investigation and engineering and management team effort involved in the recovery process. The significant time spent to recover from each safing culminates into a significant impact on overall spacecraft rates of operability.

Each safing event requires a coordinated and methodical response to recover the spacecraft. Spacecraft teams work to diagnose and understand the issue, ensuring that the anomaly is not persistent and the planned operations can be resumed. The cumulative impacts of safing events are realized when the vehicle is in flight. To manage this, margins for operational outages are prepared throughout the development phases when designing trajectories, planning critical events, and developing science campaigns. To better quantify these

¹ Galileo has the first recorded (and located) event on November 28, 1989.
978-1-5386-2014-4/18/\$31.00 ©2018 IEEE

margins on future mission concepts, a statistical approach is applied to safing data from past and present missions to model event occurrences and recovery durations.

Motivation

Initial research into modeling safing events was funded by the proposed Asteroid Robotic Redirect Mission (ARRM). ARRM's concept baselined a 40 KW solar electric propulsion (EP) system to reach asteroid 2008 EV5. [1] Low-thrust EP technologies are uniquely enabling compared to chemical propulsion systems because they allow the flight system to achieve large changes in spacecraft velocity on a comparatively small propellant mass. JPL has flown EP technologies on the Deep Space 1 (DS1) and Dawn missions, and has baselined their use on the planned Psyche mission and the Next Mars Orbiter (NeMO) mission concept. [2] [3]

EP engines produce low thrust levels, requiring near-continuous periods of operability for maneuvers. Trajectories may require weeks or months of thrusting at a time, increasing the likelihood that a safing anomaly will interrupt the planned maneuver. For example, Dawn arrived at Ceres 26 days late from a four day safing event and thrust outage. [4] Similar to DS1 and Dawn, ARRM trajectory designers performed 'missed thrust analyses' to build robustness into maneuvers. This robustness is essential due to the non-linearity of low thrust trajectories. These analyses model the impact of anomalies by injecting multi-day thrust outages into the planned trajectory and recalculating new solutions. While the missed thrust analyses give confidence a vehicle will reach the final destination, the flight system may realize lower system-level margins, such as requiring more propellant or increasing the time-of-flight. [5] [6]

EP missions are not the only architectures sensitive to safing events. Time-constrained missions, such as Europa Clipper, could be impacted by missed maneuver opportunities and brief science windows. Clipper will complete up to 45 flybys of Europa on a 14.2 day orbit cadence. The mission uses a chemical propulsion system to adjust the trajectory as the spacecraft travels within 25 km of the icy moon's surface. A safe mode entry could cause missed maneuvers, introducing a non-zero probability that Clipper could impact Europa, require a redesign of its trajectory, or miss science observations over a region of the icy moon that may not be revisited within the mission's lifetime. [7]

Missed thrust analyses on DS1 and Dawn were done using engineering best estimates for periods of inoperability. [6] To improve on this, ARRM sought to leverage nearly thirty years of beyond-Earth flight data. However, no comprehensive database, literature, or analysis of events was located. Thus began the task to create and explore this database.

2. SAFE MODE EVENT DATABASE

The safe mode event database captures nearly 240 safe mode entries from 21 beyond-Earth missions, summarized in Table

1. 196 of the records are from anomalous events throughout a mission's primary and extended phases. 12 of these records are 'cascading' events, where the spacecraft re-entered safe mode before recovery from the previous event was complete. Of these 196, 151 of the events have sufficient details and context to reconstruct the diagnosis and recovery timeline. Ancillary data is also recorded for each event including mission phase, vehicle location, and anecdotal information from the recovery process. Root causes, as specified and recorded by each mission team, are captured and binned into software, hardware, operations, space environments, or unknown categories. Unknown events are anomalies where root cause is undeterminable by the mission team.

The event database has been collected through a collaboration between JPL, NASA's Goddard Space Flight Center, NASA's Ames Research Center, and the Johns Hopkins University Applied Physics Lab – please see the Acknowledgements section. Research is ongoing to capture additional details and to locate records for missions not yet included in the database.

Table 1. The safe mode database captures events from 168 years of cumulative flight time on missions throughout the solar system.

Mission	Launch	Destination
Galileo	1989	Jupiter
Mars Global Surveyor	1996	Mars
Cassini	1997	Saturn
Deep Space 1	1998	9969 Braille
Mars Climate Orbiter ²	1998	Mars
Mars Polar Lander ²	1999	Mars
Stardust	1999	81P/Wild
Genesis	2001	Earth Sun L1
Mars Odyssey	2001	Mars
Spitzer (Telescope)	2003	Earth Trailing
Deep Impact	2005	Tempel 1
Mars Reconnaissance Orbiter	2005	Mars
New Horizons	2006	Pluto
Dawn	2007	Vesta, Ceres
Phoenix ²	2007	Mars
Kepler (Telescope)	2009	Earth Trailing
Lunar Reconnaissance Orbiter	2009	Moon
Juno	2011	Jupiter
Mars Science Laboratory ²	2011	Mars
Mars Atmosphere and Volatile Evolution	2013	Mars
OSIRIS-REx	2016	Bennu

² Cruise phase of mission only

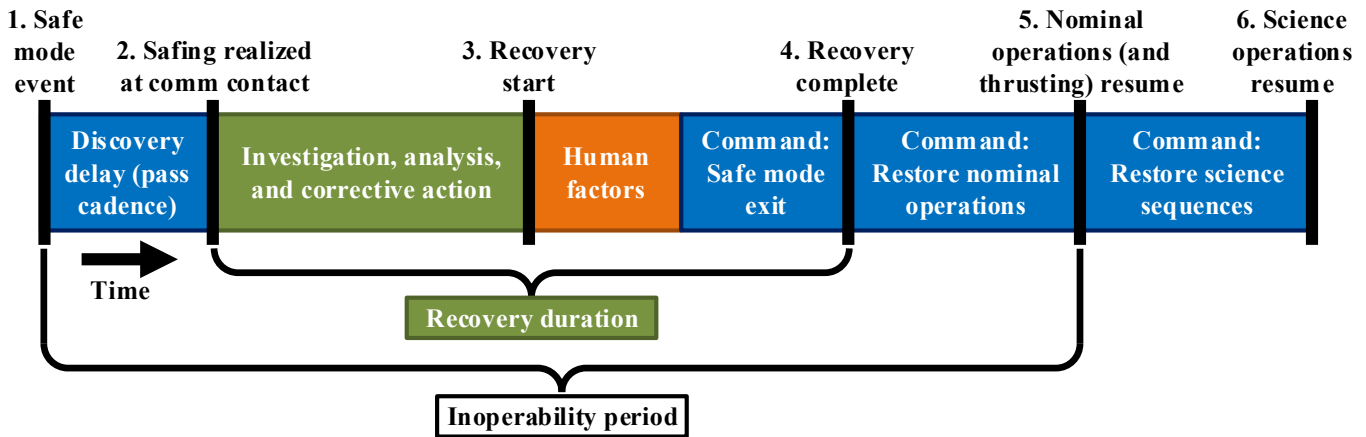


Figure 1. The safing event timeline is the framework for capturing anomaly details. The black lines are milestones and the colored blocks are periods of action. Blue periods are project specific and are quantifiable, testable, and can be levied through requirements on future missions. Green periods are statistical, derived from this research. The orange period encompasses anecdotal factors, such as days off, risk posture, recovery urgency, comm pass schedules, etc.

Event Recovery Timeline

The severity of a safing event is realized as an operational impact, requiring time to diagnose and recover the spacecraft before resuming the nominal mission sequences. To capture these details, a standardized timeline structure provides homogeneity across different missions. Figure 1 identifies the six milestones recorded for each event in the database:

1. *Safe mode event* – Safe mode entry asserted onboard
2. *Safing realized* – Ground operators receive the first indication of safe mode entry during a scheduled communications pass
3. *Recovery start* – Start of command execution to restore functionality as part of the planned safe mode recovery procedures (subjective milestone, defined by each mission)
4. *Recovery complete* – Spacecraft restored to a nominally-operating state, releasing all safe mode constraints (subjective milestone, defined by each mission)
5. *Nominal operations resume* – Restart nominal mission sequences and begin thrusting (if applicable)
6. *Science operations resume* – Science sequences resume after restoring instrument functionality (if applicable)

The timeline is agnostic to each spacecraft’s location and does not immediately reflect the impact of round trip light time (RTLTL) on the timeline periods. RTLTL can be on the order of hours, and accumulates as a mission executes required round-trip command sessions through the recovery process. Post processing can estimate the impact of RTLTL using the mission’s ephemeris from the day of the event.

Definitions

The timeline in Figure 1 is used to define three important periods used in the modeling and simulation of safing events:

Time between events is the elapsed time between two safe mode event occurrences. Since the flight durations of the missions in the database range from months (Mars Science Laboratory, etc.) to nearly two decades (Cassini), elapsed time provides a more consistent and applicable metric than events per year.

Recovery duration – highlighted in the green box below the timeline – is the elapsed time it takes the mission team to diagnose and complete initial recovery. This period encompasses both subjective and objective recovery periods. When sharing their mission’s safing data, many teams indicate that there was “no rush” or that recovery “could have been faster if needed”. Only a handful of these anecdotal notes are quantitatively recorded in the database research. As a result, the modeled recovery duration may overestimate the recovery time if a future team needs to recover quickly.

The **inoperability period** – also highlighted below the timeline – captures the overall time impact of a safing event. The spacecraft is defined as inoperable when not executing the planned mission sequences or is inhibited from thrusting. Outside of this period, the spacecraft is considered operable and maneuverable. Summing all of the inoperability periods over a mission’s full duration gives the net **inoperability rate**, so the mission’s **operability rate** is therefore 1 minus the inoperability rate.

3. MODELING OCCURRENCES AND RECOVERIES

Initial data exploration and model development techniques have been applied to the event database to characterize top-level behaviors of the data and identify candidate models for further investigation. After fitting and validating these models, the selected model becomes part of a tool that enables portability of the mission-specific data to future architectures. Both time between events and recovery durations are investigated in this effort. Different statistical distributions are fit to the empirical data to look for trends

and represent the raw data in meaningful ways. The initial modeling is based on five foundational assumptions:

- *All missions are equal and identically distributed* – Class, cost, launch year, destination, complexity, instruments, propulsion type, and other factors are not evaluated.
- *All events are in the same population* – Data is lumped into common datasets, meaning different numbers of events between missions are not considered. All events follow the same natural process.
- *All events are random and independent* – Variations in elapsed time-of-flight variations are not considered (events that might be more common early in the mission) and past events have no influence on future events.
- *All events are equal* – Influence of root cause, system complexity, and other factors are not considered in recovery durations.
- *All recoveries are perfect* - Upon recovery, the spacecraft has the same probability of safing again as it did prior to the event.

These assumptions provide a starting point for the initial investigation and model development. Lumping the time between events and recovery duration data points into independent and identically distributed populations creates large datasets that can be assessed for statistical significance. However, the appropriateness of these assumptions needs to be verified and potentially confounding variables will require further analysis. The Recommendations and Forward Work section discusses future efforts to test these assumptions and understand their influence on the results.

Developing the Datasets

Datasets for analysis and modeling are assembled directly from the safing event database with no manipulation unless explicitly noted. Following the assumption that all missions and events are equal and from the same sample population, all valid data points from all missions are combined into single, large datasets. Though the motivation of the research comes from missed thrust analyses, the initial inspection shows no significant observed difference in behavior of safing events on the two EP missions (Dawn and DS1) versus the other missions in the database. Histograms of the time between events and recovery duration datasets are in Figure 2, overlaid with best-fits (discussed in the next subsection).

The time between events dataset is first organized into mission-specific subsets by calculating the elapsed time between unplanned safing events within a mission's flight duration. This methodology does not consider the impact of other factors, such as root cause, occurrences of planned safe mode entries, or if the mission transitions between phases. This process treats cascading events (discussed previously) as a single data point since nearly all of the cascading events happen within a day of the first event. Elapsed time from launch to the first event is included as a valid elapsed time. Once each mission's time between events subset is calculated, records from all missions are combined to assemble the full time between events dataset.

The recovery duration dataset is assembled directly from the recovery timelines associated with each safing event. Similar to time between events, this process ignores the impact of other factors and treats cascading events as a one duration. Some specific recovery duration data points are excluded; One multi-month safing event and all Galileo recovery duration data is not evaluated per specific recommendations from team members involved on those projects.




Figure 2. Histograms of both datasets are shown with fitted Weibull distributions. The longest time between events record is about 4.3 years (left), while no included event had a recovery duration longer than 10 days (right).

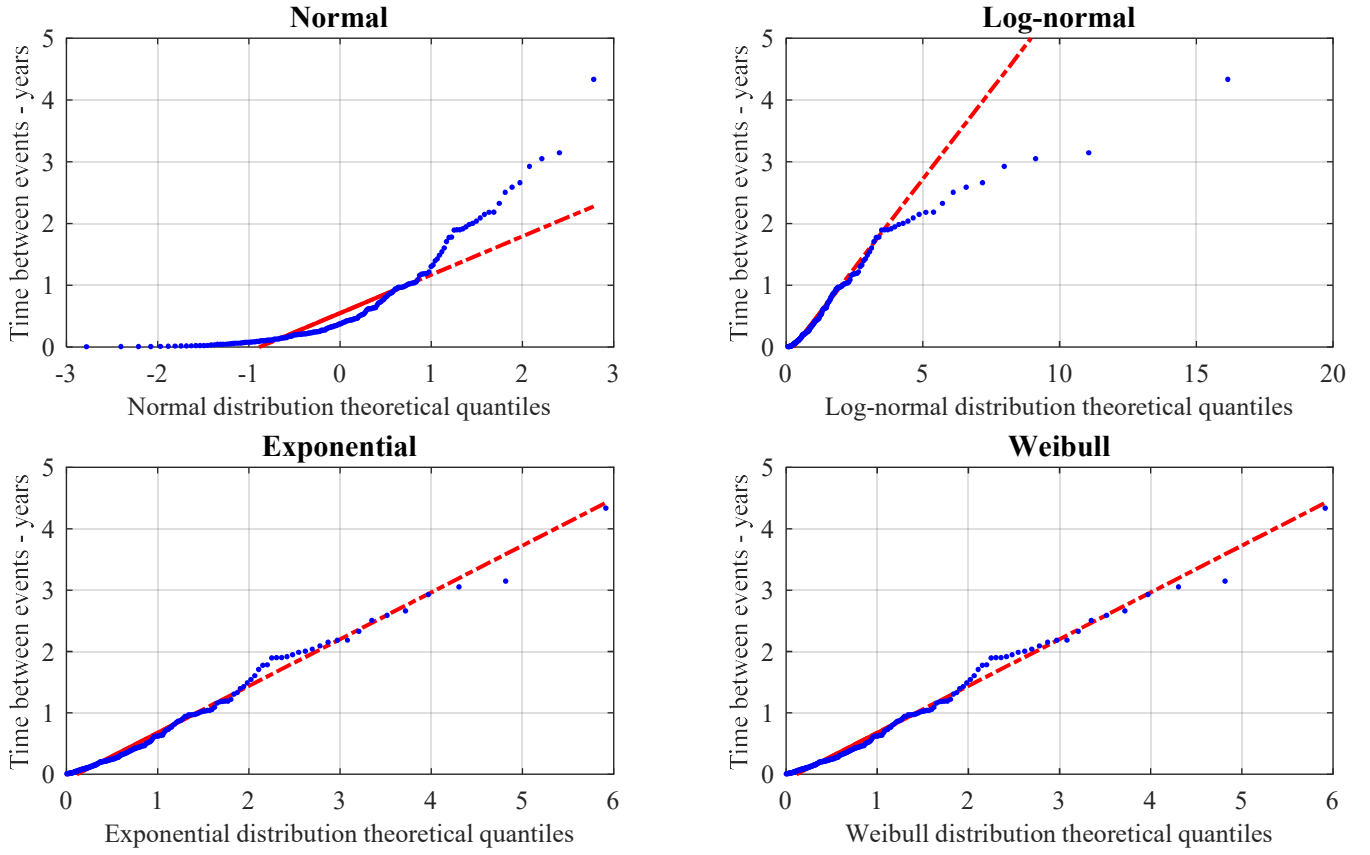


Figure 3. Quantile-quantile plots visually compare the time between events dataset (y-axes) against the theoretical quantiles of four distributions (x-axes). Closeness-of-fit between the empirical data (blue dots) and the theoretical distribution (straight red line) is inspected. The same process is also done for the recovery duration dataset.

Data Modeling and Distribution Fitting

To model the datasets in Figure 2, several different statistical distributions were applied and compared against the empirical data. Using the histograms as a starting point, both datasets were compared to the normal, log-normal, exponential, and Weibull distributions. This evaluation was done using the quantile-quantile (q-q) plots shown in Figure 3. The q-q plot is a technique where two distributions are compared graphically to investigate whether they come from similar populations. If two distributions come from the same population, then the two datasets should closely overlap. Both the time between events and recovery duration datasets (blue dots) were compared to theoretical distributions (straight red lines), as illustrated in the four sub plots of Figure 3. Both the normal and lognormal distributions fit parts of the data, but quickly diverge. The exponential and Weibull distributions have excellent closeness-of-fit for the full datasets, with the Weibull ultimately providing a slightly better representation.

The equation of the Weibull probability distribution function (as plotted in Figure 2) is given in Equation 1, where k is the shape parameter and λ is the scale parameter.

$$f(x; k, \lambda) = \begin{cases} \frac{k}{\lambda} \left(\frac{x}{\lambda}\right)^{k-1} e^{-(x/\lambda)^k} & x \geq 0 \\ 0 & x < 0 \end{cases} \quad (1)$$

The Weibull is commonly used in reliability and failure engineering and is fully defined by the scale and shape parameters. The scale defines the magnitude of the horizontal axis value, while the shape gives insight into the behavior of the sample population. The similarities between the exponential and Weibull seen in Figure 3 are not surprising; When the Weibull shape parameter is 1, it mathematically reduces to the exponential distribution. This extra parameter gives the Weibull an additional degree of freedom to better fit the data. When the Weibull shape is 1 (exponential), it indicates that the event rate does not change as a function of time. A shape less than 1 indicates a decreasing event rate over time, while a shape greater than 1 highlights an increasing rate.

The histograms in Figure 2 are shown as cumulative distribution functions (CDFs) in Figure 4. The empirical data (which falls in the same x-axis location as the histograms) is shown in various colors to represent the contributions each mission anonymously. Since CDFs show cumulative probability, the empirical data is distributed evenly on the y-axis and is overlaid with the integral of the Weibull fits. The time between events CDF shows there is a 90% probability the next safing will occur within ~1.6 years of the previous event. Similarly, there is a 90% probability that the recovery duration of any event will be ~5.1 days or fewer.





Figure 4. The Weibull fits for the *time between events* CDF (top) and the *recovery duration* CDF (bottom) closely match the empirical data. The colors anonymously highlight the contribution of each mission to both datasets, with totals shown in parentheses. 5 missions have no locatable recovery duration data.

The Weibull parameters for the time between events and recovery duration datasets are found using MATLAB's Weibull maximum likelihood estimator (*wblfit*) and are shown in Table 2.

Table 2. Weibull parameters fit to the empirical data

Distribution	Scale	Shape
Time between events (years)	0.62	0.87
Recovery duration (days)	2.46	1.14

Though the Weibull was selected due to the closeness-of-fit to the empirical data, the calculated parameters provide initial insight. Both distributions have a shape relatively close to 1, highlighting a near constant occurrence rate and confirming the similarities to the exponential distribution seen in Figure 3. The time between events shape of 0.87 indicates a slight decrease in the day-to-day likelihood of safing again. This could be explained by team operations and procedures becoming more routine and robust over time. The recovery duration shape of 1.14 highlights an increasing recovery rate each day after event entry. This may be due to initial difficulty in securing communications pass time immediately after discovery, and could represent the process of mission teams gathering data and working to better understand the anomaly over time.

The most significant impact of the calculated Weibull fits is the portability of the data for future use. With only the shape and scale parameters, a pseudo-random number generator can be used to return Weibull-distributed times between safing events and recovery times. This allows the safing data to be leveraged in simulations and mission design independent of the raw datasets. While the Weibull parameters will change as more missions are included, deviations from the parameters presented here should be relatively minor due to the comparatively large sample size of both datasets.

4. APPLICABILITY TO FUTURE MISSIONS

Future missions can combine Weibull parameters with mission specific details to begin to quantify the cumulative impact of safing events on their architectures. This process addresses architecture enabling assumptions during mission formulation, such as depending on a minimum operability rate or planning for finite time at a destination. To illustrate the potential use in mission design, a simulation methodology helps to quantify the cumulative impact of safing events over a mission's life. When driven by a Monte Carlo front-end, the outputs of each simulation are used to generate a probability distribution of the Minimum Operability Rate (MOR). Analyzing the MOR distribution allows mission designers to compare risk to performance, illustrating likelihood of achieving an overall operability rate or better over the mission's full duration. Findings can also be valuable to probabilistic risk assessments of spacecraft reliability and performance.

A flowchart of the simulation methodology is shown in Figure 5. The simulation makes use of a Weibull-based pseudo-random number generator (*wblrnd* in MATLAB) based on the distribution parameters found in the analysis. Both the first and subsequent safing occurrences are determined using the time between events Weibull fit. When an event occurs, the inoperability period is found by summing the discovery delay, recovery duration from the Weibull fit, and the time to restore nominal operations. The discovery delay block accounts for pass length, where an anomaly may be realized near-instantly if it occurs during a scheduled pass. After totaling the inoperability period, the simulation then predicts the next safing occurrence and checks if the next entry is considered a cascading event. When a cascading entry occurs, the simulation adds the previous, interrupted recovery duration to a re-simulated inoperability period for the cascading safing entry. Multiple cascades are possible. In the simulation, the cascading event occurrences are recorded as single events and single, potentially longer recovery

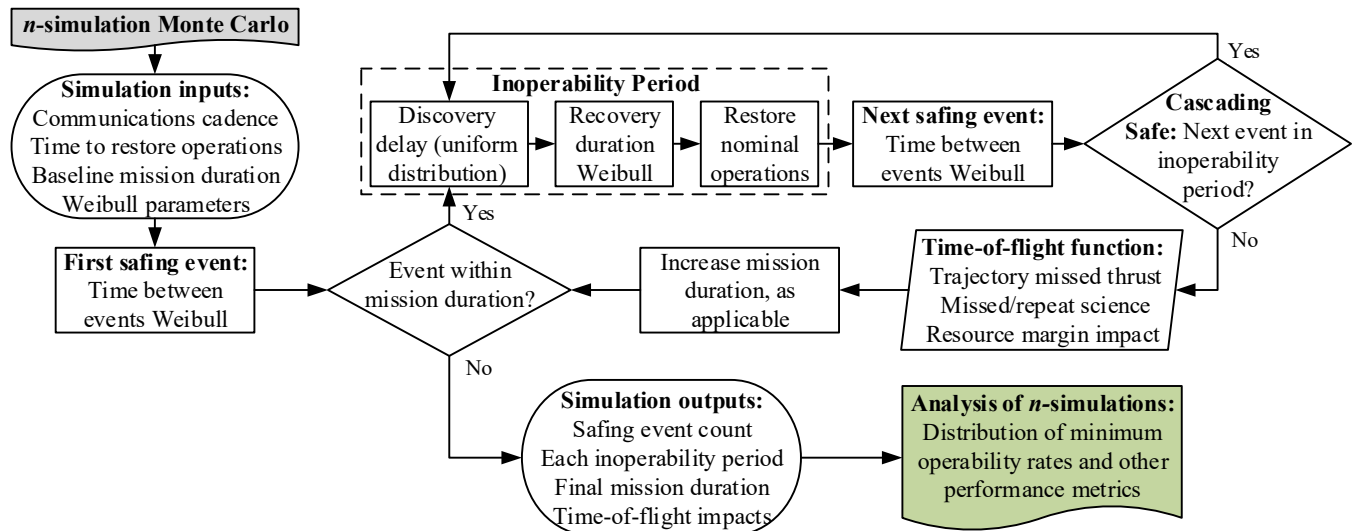


Figure 5. The simulation process enables performance modeling of a candidate mission architecture. This methodology combines the Weibull fits with mission-specific details to investigate minimum operability rates.

durations in the outputs. Metrics and impacts of cascades are noted separately in the outputs.

Following a single event's inoperability period, the simulation enters a time-of-flight function. This interface allows external tools, such as trajectory optimization software, to quantify impacts to the planned mission duration. For low-thrust mission architectures, the cumulative additional time-of-flight from the safing events may be significant enough that additional events could occur. The simulation will continue until the next safing event occurs at a point in time beyond each simulation run's full mission duration.

Simulating a Candidate Mission

A mission with a continuous multi-year thrusting phase is simulated using a Monte Carlo model in MATLAB. The inputs used in the simulations are shown in Table 3. Since no specific trajectory is under consideration, a 4 hour pass every three days and a 1:1 time-of-flight increase are baselined.

Table 3. Simulation parameters for the sample mission

Metric	Value
Baseline mission duration	6 years
Discovery delay / Pass cadence	Pass every 3 days
Pass length	4 hours
Time to restore nominal operations	12 hours
Time-of-flight function increase, relative to inoperability period	1:1

One million simulations were performed in the model, totaling around 9 million safing events shown in Figure 6. With the 1:1 time-of-flight function, time-of-flights increased up to 6.24 years (99.7th percentile). The left histogram indicates that between 0 and 20 events (99.7th percentile) are

possible during the mission duration, while 6 – 11 events are most probable. The right histogram shows the range of simulated inoperable periods from each safing entry. No durations are less than 12 hours, with the peak around 3.5 days; this is driven by the simulation inputs for discovery delay and command time to restore nominal operations. The right tail is driven by the recovery duration Weibull fit, decaying to an inoperable period of 14 days (99.7th percentile).

Over the Monte Carlo runs, the cumulative inoperability period of the simulations ranged from 0 to 88 days (99.7th percentile). The overall operability rate distribution is shown as a survivor function in Figure 7. For this simulation, the distribution shows all simulation runs converged above a MOR of 94%, descending to an expected near-zero likelihood of achieving a 100% operability rate. The 99.7th percentile, marked by the green star, maps to a 96.1% MOR. This is equivalent to at least 350.8 days of operability or no more than 14.2 days of inoperability per year. The 95th percentile lies at a 97.1% MOR, about 10.6 days of inoperability per year.

While the operability rates presented here will change as investigations apply mission-specific constraints, these simulations begin to bound the expected operability rates of beyond-Earth spacecraft. Coupled with the Weibull fits, this simulation methodology provides valuable insight into plausible vehicle operability rates early in the formulation process. A design team is able to target a desired operability rate, quantify the likelihood of achieving that rate, and compare the results against project risk posture and performance margins.





Figure 6. The distribution of total mission safings (left) peaks at 8 to 9 events, while the tail of the distribution of the total inoperable period of each safing (right) has recovery durations that extend beyond 30 days due to cascading events.




Figure 7. The shape of the distribution is a convolution of the simulation inputs with the Weibull distributions. 94.2% is the lowest observed MOR, and 0.08% of the simulations had a 100% operability rate.

Sensitivity Investigation

To investigate sensitivities in MOR, the Monte Carlo model is used to look at mission simulations with durations between one and ten years and with a discovery delay ranging from one pass per day to one pass per week. The time to restore nominal operations is kept constant at 12 hours per event, the pass length is 4 hours, and the time-of-flight function remains at a 1:1 duration increase. The results are shown in Figure 8.

The sensitivity investigation shows that MORs are plausible from 89% to 97% (99.7th percentile) depending on mission duration and communications cadence. Decreasing the discovery delay with more frequent passes results in notable,

multiple percent increases in MOR. More generally, the sensitivity analysis shows a driving relationship between MOR and baseline mission duration. Shorter missions have comparatively lower MORs when evaluated at these high percentiles. While shorter missions will have a higher percentage of simulations with no safing events, the 99.7th and 95th percentile MORs of shorter missions are dominated by simulation runs with many safing events. In these situations, the simulation duration is sufficiently short such that long periods between events will fall outside the simulation's mission duration and are unable to 'boost' the MOR percentiles. Finally, increasing the time-of-flight function beyond a 1:1 impact results in slightly higher MORs, since the baseline mission duration is effectively increased. However, there may be other impacts to vehicle margin and resources that would be quantified through other tools.




Figure 8. MOR is sensitive to baseline mission duration, with the shortest missions expecting a MOR above 89%.

5. RECOMMENDATIONS AND FORWARD WORK

The analysis, modeling, and sensitivity results lead to initial recommendations in managing inoperability period margin on future mission architectures. Previous publications on managing margins for low-thrust missions have highlighted the sensitivities between cost, propellant mass, flight time, and launch vehicle selection. [5] However, assumptions for missed thrust periods have traditionally been modeled as overall system performance rates (such as 95% duty cycle on Dawn) because the tools and data for deriving quantitative insights were not yet available. [6] With the foundation of the safing event database, future mission teams can use the Weibull fits and simulation methodology to develop confidence in minimum expected operability rates of their own architectures.

Overall, this body of work provides key takeaways that can begin to bound the behavior and impact of safing events on future mission architectures:

- Safing events are plausible on missions of all types, classes, destinations, and architectures with no observed difference between electric and impulsive propulsion missions.
- The mean of the time between events Weibull fit is about three events every two years. No included mission has recorded a recovery duration longer than 10 days, with 90% of the records under ~ 5.1 days.
- Exploring MOR distributions through simulation establishes a tailorable methodology to understand a vehicle's performance based on the Weibull fits.
- MOR is sensitive to communications pass cadences on missions of all lengths, allowing performance margin to be bought when it is needed

Further Investigations

The analysis and findings presented here are the first exploration of the records captured in the safing event database. Both the Weibull modeling and simulation techniques will continue to undergo a validation process to ensure that the outputs are appropriate, accurate, and best represent the collected mission data.

Processes to improve the interpretation and inclusion of all of the event records are ongoing. Recovery durations may decrease slightly when the full impact of round trip light time is measured and factored in. However, insufficient data has been located on the number of round-trip command sessions used on past missions. Techniques are being developed to represent missions with no safing events (such as GRAIL) and to include events that were explicitly excluded per recommendations from those mission teams. The impact of statistical censoring on the datasets has not been fully evaluated, which will factor in the time between a completed mission's last event and end-of-mission, or represent current operable periods on ongoing missions that may or may not have a future event.

Another primary focus will revisit the initial assumptions to test the value of developing methodologies to improve the representations of the raw dataset. While the assumptions that all data is identically distributed and independent are simplifications of each mission, all missions studied here were designed to survive and operate as beyond-Earth spacecraft. Though their top-level objectives are diverse, these vehicles share many common architectural elements. Thus, developing and validating new conclusions may not reveal significantly deeper insight than conclusions derived under the initial assumptions.

Analysis of the datasets are expanding beyond the lumped Weibull fits. Binning methods are used to create separate

probabilistic distributions within the datasets, such as comparing trends between primary and extended missions, cruise and orbit, etc. Exploring variations in ancillary factors such as destination, cost, time-of-flight, and root cause provides a first look into trends and differences. Beyond binning, statistical techniques will be applied to measure the influence and relationships between missions in each dataset. Confidence intervals can be applied to the Weibull distributions, and mixed Weibull distributions could improve the closeness-of-fit to the outliers seen in the q-q plots. Hypothesis tests such as the Kolmogorov-Smirnov, t-test, or chi-square test can look for statistical connections within the datasets. These questions are the beginning of understanding the 'weight' of each mission on the resulting distribution fits. While binning creates an absolute sorting, weighting techniques could be applied for more tailored simulations. For example, the NeMO mission concept could consider Mars orbiter data more significantly than others.

The simulation methodology will evolve with database and dataset improvements. Integration of trajectory software, such as MALTO, through the time-of-flight function is a major focus. While vectors of time between events and recovery duration values can be generated and exported to existing tools, closed-loop interfacing with trajectory optimization tools can fully capture the impacts of discovery delays, round trip light time, and cascading safes. While the inclusion of trajectory optimizers will slow down the speed of the Monte Carlo, the simulations are easily parallelized for high-performance computing environments.

Finally, the event database will continue to grow as new missions are added and missing records are located. While the dates of safing event entries are relatively straightforward to locate through records and interviews, only half of the missions have recorded recovery durations for all of their safing events. Beyond this, five of the missions have no recovery duration for any events, and nearly every safing record is missing various details related to the safing event timeline. While different mission practices, lost records, and retired personnel mean that some of this data will never be locatable, the authors welcome assistance in verifying and growing the database.

6. CONCLUSIONS

Safe mode events are the predominant contributor to periods of spacecraft inoperability. Safings consume mission resources to diagnose, recover, and prevent further vehicle faults. Exciting mission architectures, such as the NeMO concept, the planned Psyche mission, and Europa Clipper are sensitive to the cumulative impact of safing events on the spacecraft's planned minimum operability rate. To better understand and quantify this risk, a comprehensive database of spacecraft safe mode events from past and present missions has been collected, stemming from a collaboration between JPL, NASA's Ames Research Center, NASA's Goddard Flight Center, and the Johns Hopkins University Applied Physics Laboratory.

Initial exploration and analysis of the events in the safing database highlights novel research. A standardized timeline is introduced to provide homogeneity in collecting detailed records between missions. Several initial assumptions are applied, allowing data from all missions to be combined in two large datasets for elapsed time between events and recovery durations. Various statistical fits are tested against the empirical data, with Weibull distributions providing the best closeness-of fit. The shape and scale parameters of the two Weibull distributions are discussed and presented to enable portability of the datasets beyond the raw data.

These parameters and other mission-specific details are applied in simulations driven by a Monte Carlo model. A candidate six year mission with realistic operations variables shows a minimum operability rate (MOR) of 96.1% is plausible at the 99.7th percentile. A brief sensitivity investigation highlights the range of plausible MORs by varying the baseline mission duration and pass cadences. The investigation shows MORs up to 97% (99.7th percentile) are achievable through realistic optimizations of the spacecraft design and operations plan. In future applications, trajectory optimization and other tools can be interfaced through simulation's time-of-flight function to develop high-fidelity, mission-tailored MOR distributions.

Some exploratory work remains to improve interpretation of records in the database. Future investigations will improve the modeling and simulation by revisiting the initial assumptions and applying different statistical techniques to validate the work presented here. This includes testing hypotheses concerning event and recovery behaviors, uncovering potential sub-groupings in the event database, and exploring alternative modeling approaches.

Finally, this paper is the first of many investigations resulting from the records in the safing event database. Improved understanding of the occurrence and severity of safing events will be valuable to future missions, especially those enabled by electric propulsion architectures or time-constrained operations. Recommendations are presented to guide future mission architects, trajectory designers, and spacecraft operators as they model and predict the performance of their systems.

ACKNOWLEDGEMENTS

The authors thank Kelli McCoy and Steve Franklin for starting the safing event database. This research was made possible due to the collaboration and assistance from the following individuals. At the **Jet Propulsion Laboratory**: Dawn – Paul Fieseler, Marc Rayman, Dan Grebow, Bruce Waggoner; MRO – Dan Johnston; Cassini – Laura Burke, Julie Webster; Mars Odyssey – Bruce McLaughlin, David Lehman; Juno – Tracy Drain; Spitzer – Eric Rice, Joseph Hunt; Phoenix – Grant Faris; Deep Space 1 – Tracy Neilson; Genesis and Stardust – Ed Hirst; MSL – David Oh; Galileo – Robert Gounley. Records were also located through JPL's incident database, enabled by Ian Colwell's analytics tool.

Johns Hopkins University Applied Physics Laboratory: New Horizons – Chris Hersman. **NASA Ames Research Center**: Kepler – Charlie Sobeck, Marcie Smith. **NASA Goddard Space Flight Center**: OSIRIS-REx and LRO – Dave Everett; MAVEN – John Nagy, Martin Houghton, Rich Burns.

This work was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

REFERENCES

- [1] NASA's Asteroid Redirect Mission Concept Overview, <https://hou.usra.edu/meetings/lpsc2016/eposter/2217.pdf>
- [2] NASA Selects Five Mars Orbiter Concept Studies, July 2016. <https://www.nasa.gov/press-release/nasa-selects-five-mars-orbiter-concept-studies>
- [3] Psyche, Mission to a Metal World, <https://jpl.nasa.gov/missions/psyche/>
- [4] D. Grebow, G Whiffen, D. Han, and B. Kennedy, *Dawn Safing Approach to Ceres Redesign*, AIAA/AAS Astrodynamics Specialist Conference, Sept. 2016.
- [5] D. Landau, T. Kowalkowski, J. Sims, T. Randolph, P. Timmerman, J. Chase, and D. Oh, *Electric Propulsion System Selection Process for Interplanetary Missions*, AIAA/AAU Astrodynamics Specialist Conference, Aug. 2008.
- [6] M. Rayman, T. Fraschetti, C. Raymond, and C. Russell, *Coupling of system resource margins through the use of electric propulsion: Implications in preparing for the Dawn mission to Ceres and Vesta*, Acta Astronautica, Jan. 2007.
- [7] T. Bayer, B. Buffington, J.F. Castet, M. Jackson, G. Lee, K. Lewis, J. Kastner, K. Schimmels, and K. Kirby, *Europa Mission Update: Beyond Payload Selection*, 2017 IEEE Aerospace Conference, March 2017.

BIOGRAPHY



Travis Imken received a M.S. in Aerospace Engineering from the University of Texas at Austin in 2014. He is in the Project Systems Engineering and Formulation Section at JPL, supporting the project systems team on the 2018 InSight Mars Lander and serving as the project systems engineer on the RainCube “Ka-band radar in a CubeSat” mission. Prior to this, Travis was involved with the ARRM mission and the Lunar Flashlight and NEA Scout deep space CubeSats. He is also involved with JPL's Innovation

Foundry, serving as a systems engineer on Team X/Xc as well as a small satellite expert with the A Team.



Thomas Randolph received his B.S. in Aerospace Engineering from the University of Southern California and his M.S. in Mechanical and Aerospace Engineering from Princeton University while doing research in ion and magnetoplasmadynamic thrusters. After leaving school in 1992, Tom worked at Space Systems Loral where he was initially the Cognizant Engineer for the successful western qualification of the Russian SPT-100 Hall thruster and later the Electric Propulsion Product Manager for the first commercial flights of a Hall thruster on MBSAT. After starting work at JPL in 2003, Tom has been the Product Delivery Manager for the ST7 colloid microthruster system, the Project System Engineer on the Low Density Supersonic Decelerator Technology Demonstration Mission, the Technical Group Supervisor of the Project Systems Engineering Group, the Project Systems Engineer on the Asteroid Redirect Robotic Mission, and is currently the Project Manager for ASPIRE, Advanced Supersonic Parachute Inflation Research Experiments.



Michael DiNicola is a senior engineer at the Jet Propulsion Laboratory in the Systems Modeling, Analysis & Architectures Group. In his ten years at JPL he has been part of several proposals, risk analyses and model development efforts. He currently works on the Europa Clipper Project developing probabilistic models to assess Planetary Protection and science requirements. He also provides statistical analysis as part of the NASA Instrument Cost Model Team. Michael attained his B.S. in Mathematics from UCLA and M.A. in Mathematics from UCSD.



Austin Nicholas received a B.S. degree in Aerospace Engineering from University of Illinois at Urbana-Champaign in 2011 and an M.S. in Aeronautics and Astronautics from the Massachusetts Institute of Technology in 2013. He currently works for JPL in the Project Systems Engineering & Formulation Section. His primary assignment is in the Mars Program Formulation Office developing architectures and vehicle concepts for Mars Sample Return. Recent work has included solar electric propulsion spacecraft with co-optimized flight systems and trajectory, developed for the Next Mars Orbiter concept. At MIT, he worked on attitude and cluster control for a formation-flight Cubesat mission using electrospray propulsion and architecture exploration for low-cost human missions to the lunar surface.